

Appendix 1: Glossary of Terms

Action

An action is a conceptual memory of a procedure the agent can perform, e.g., grab a cup, which can have several associated algorithms for execution in Sensory-Motor Memory.

Activation

Activation can represent different things, but, generally, it represents the salience or usefulness of an element. Nodes, Links, and other elements such as coalitions, codelets, and behaviors, have activation.

Aggregate-Coalition-Activation Trigger

A broadcast trigger that fires when the sum of the activations of all coalitions in the Global Workspace is greater than a certain threshold.

Attention Codelet

An Attention Codelet is a type of task that tries to bring particular Workspace content in the Current Situational Model to the Global Workspace. Upon finding such content it creates a coalition containing the content and adds it to the Global Workspace to compete for consciousness.

Autonomous agent

A system situated within and a part of an environment that senses that environment and acts on it, over time, in pursuit of its own agenda, and so as to affect what it senses in the future. Most animals and software agents are examples.

Background Task

A task, typically internal to a module, which implements a process integral to the module's functionality. Such tasks run repeatedly “in the background”. An example is the `SensoryMemoryBackgroundTask` that runs the agent's sensors each time it is executed.

Base-Level Activation

Activation that represents the general usefulness and relevance of an element in the past. Its value can range from 0.0 to 1.0 inclusive and is modified by learning.

Behavior

A behavior is an instantiated scheme. Schemes are instantiated in Procedural Memory. Instantiation binds variables in the scheme's context and result, based on the current conscious

broadcast. Instantiated schemes (behaviors) are then sent to Action Selection where they compete with each other for selection and execution.

Behavior Network

An implementation of the Action Selection module inspired by Maes' (1989) Behavior Net. It selects a behavior to execute at the conclusion of each cognitive cycle. This selection is based on the activation of the Behavior, which is itself dependent on the current contents of consciousness as well as the relationship the Behavior has with other Behaviors currently in the Behavior Network.

Bit Vector

Bit vectors are typically large vectors with values of either 0 or 1. They are used in the Sparse Distributed implementation of Episodic Memory as a means to store episodic memories.

Broadcast Trigger

A process that monitors the Global Workspace for a particular condition. Upon detecting that condition it prompts the Global Workspace to hold a competition for consciousness.

Coalition

A coalition contains a portion of the Current Situational Model that is brought to the Global Workspace as a unit along with the attention codelet that created it. Coalitions compete for consciousness based on their activation derived from both the activation of the content and of the attention codelet. The winning coalition's contents are broadcast throughout the system.

Cognitive Cycle

In the LIDA Model, a sequence of processes that occur over time beginning with the sensing of stimuli, proceeding with understanding, then continuing with attention (consciousness), and ending with action selection and execution. Thus this cycle can be conceived of as having three main phases: understanding, attention (consciousness), and acting. Multiple cognitive cycles may be occurring in parallel (cascade): as one cognitive cycle has reached the action phase another could be in the attention phase while a third is in the understanding phase. A cognitive cycle can be thought of as a cognitive "atom" — the basic unit on which higher-level cognitive processes are built. The agent's "life" can be viewed as consisting of a continual cascade of these cognitive cycles.

Complex Link

A Link that has a Node for its source and a simple Link for its sink. Complex Links cannot be sinks.

Conscious Contents Queue

The Conscious Contents Queue holds the contents of the last several seconds of broadcasts and allows LIDA to ground and produce time-related concepts including duration.

Current Activation

An activation that represents an element's current salience. Can take a value from 0.0 to 1.0 inclusive.

Current Situational Model

The Current Situational Model is a buffer in the Workspace containing structures representing the agent's current situation. It contains a perceptual scene comprised a real and virtual part. The former consists of the current internal representation of the real world while the latter consists of the current internal rememberings, plannings or imaginings of the agent.

Decay Strategy

A strategy that decreases an activation value over time (in ticks). May be implemented using a sigmoid function for example.

Declarative Memory

Declarative Memory is a module in the LIDA Model, also called long-term episodic memory. Declarative Memory can be subdivided into the autobiographical memories of events, e.g., your last birthday celebration, and the semantic memory of facts like "Paris is the capital of France."

ElementFactory

The ElementFactory's main purpose is to create new element objects (e.g. Nodes) in a centralized, configurable way. This class implements the singleton and factory design pattern and thus provides an encapsulation of the constructors of basic framework elements. It is the preferred way to obtain new Nodes, Links, some tasks, and Strategies.

Episodic Memory

Episodic Memory is a memory for events, the what, the where, and the when of a given episode. In the LIDA Model it is comprised of Transient Episodic Memory and Declarative Memory.

Excite Strategy

A strategy that increases an activation value by a specific amount. May be implemented using a sigmoid function for example. Note that exciting activation value with a negative amount will likely decrease the activation.

Feature Detector Algorithm

A feature detector algorithm is a specific type of task. Its function is to detect a particular feature in Sensory Memory and excite a PAM Node representing this feature in Perceptual Associative Memory.

Global Workspace

A module where coalitions, created and added by AttentionCodelets, compete for consciousness. The content carried by the winning coalition becomes the current contents of consciousness and it broadcast throughout the system.

Global Workspace Theory

A theory of consciousness and cognition proposed by Baars, which describes consciousness functionally as a process by which many unconscious processors compete for access to a “Global Workspace.” The winning processor has its contents broadcast globally throughout the cognitive system. It is one important foundation theory of LIDA, which implements the Global Workspace and global broadcast.

GUI Command

GUI-generated events, e.g. a button press, that require a response in the model (the agent) are implemented as GUI commands. These are thus encapsulations of commands from the GUI to be performed in the model. GUI commands implement the Command design pattern.

Individual-Coalition-Activation Trigger

A broadcast trigger that initiates a competition for consciousness if there exists a coalition with activation above a certain threshold in the Global Workspace.

LIDA Cognitive Model

An integrated artificial cognitive system that attempts to model a broad spectrum of cognition in biological systems, from low-level perception/action to high-level reasoning. It is empirically grounded in cognitive science and cognitive neuroscience. In addition to providing hypotheses to guide further research, the model can support control structures for software agents and robots.

LIDA Software Framework

A software framework developed to allow a customizable implementation of the LIDA Cognitive Model. Specifically, an application programming interface (API) containing Java classes and interfaces that implement modules and processes according to the computational architecture of the LIDA Model. Also, due to its high degree of flexibility, this framework is suitable for wider use in developing other cognitive architectures.

Link

A connection between a source Node and a sink, which may be either a Node or a *simple Link*. A Link has a LinkCategory, which is the Link's conceptual meaning.

LinkCategory

LinkCategory is the conceptual meaning of a Link, e.g. membership, lateral, causal, parent, none, etc.

Listener

Modules need to communicate with other modules. To implement this, we use the observer design pattern. In this pattern a "listener" module registers itself with another "producer" module. Each time the producer has something to send, it transmits the information to all of its registered listeners.

Logger

Logger is a Java utility used to log messages about important events and warnings from specific framework components.

Module

Similar to a module in the LIDA Cognitive Model, we generically define a module in the Framework. Modules contain information of various kinds, and include processes and algorithms with which to manipulate, modify, and move that information. All modules implement a generic module interface (API) but each one also has its own API that defines its particular functionality.

Node

A Node is a major unit of representation in the framework, which can represent a feature, object, category, event, etc.

NodeStructure

NodeStructure is a graph structure, containing Nodes and the Links between them. It constitutes the main internal representation for many framework modules.

No-Broadcast-Occurring Trigger

A broadcast trigger that initiates a competition in the Global Workspace when no broadcast has occurred for a certain period of ticks.

No-Coalition-Arriving trigger

A broadcast trigger that initiates a competition in the Global Workspace when no new coalition has been added to the Global Workspace for a certain period of ticks.

Perceptual Associative Memory (PAM)

A module that makes sense of the incoming sensory information, creating additional understanding. Contains feature detector processes, which detect features in Sensory Memory and excite their associated PAM Nodes. The Node-and-Link memory of PAM is implemented as a semantic net with activation passing. Nodes propagate their activation to the other Nodes they are linked to. Those Nodes that gain activation above their threshold become part of the current percept, which becomes a part of working memory (enters the Workspace). PAM is essentially recognition memory.

Procedural Memory

A module where schemes are stored. Here schemes can be activated by each conscious broadcast, and, if this puts them over their threshold, are instantiated as behaviors and sent to Action Selection. Procedural Memory essentially remembers what to do in what circumstance to achieve a given result.

Receptive Field

The set of units from which a feature detector (PAM Node) takes as its input. The units can be other PAM Nodes or elements in Sensory Memory.

Scheme

Schemes are memories, stored in Procedural Memory, for various actions (procedures), which may be selected for execution. A scheme consists of a context, an action and a result. With some reliability, a scheme's result is expected to occur when its action is taken when its context is satisfied.

Sensory Memory

A module that senses the environment and holds the incoming sensory stimuli for a brief period. Low-level processing (e.g. color, texture, segmentation in the visual modality) of the sensory stimuli occurs here. Feature detectors have their receptive fields in Sensory Memory.

Sensory-Motor Memory

A module that stores algorithms (motor plans) for actuator execution. This module receives selected Behaviors from the Behavior Network and determines the appropriate algorithm for executing the action. The agent's actuators then execute the algorithm. During the execution it receives information from Sensory Memory in an implementation of the so-called "dorsal" stream.

Simple Link

A Link that connects two Nodes.

Strategy

Strategies are implementations of the Strategy design pattern. They are encapsulations of calculations, algorithms, etc. which may have several desirable implementations. This allows for flexibility, for example, either a linear function or a quadratic function could be used as a Node's excite strategy (provided they implement the ExciteStrategy interface).

Structure-Building Codelet

A Structure-Building Codelet is a specific type of task. It operates with representations in the submodules of the Workspace making associations among representations, creating new representational structures, and combining similar ones. They play a role in maintaining the agent's model of its current situation in the Current Situational Model.

Submodules

Modules can have submodules, which are modules themselves that are nested inside another module. For example, the Workspace module in LIDA has several submodules such as the Current Situational Model.

Task

Tasks are encapsulations of small, specialized, routines (algorithms), which typically run repeatedly at a specified rate. Tasks consist of an algorithm, an execution frequency, a time of next execution, and a status. Modules or even other tasks can create tasks to perform the small operations needed to achieve their specific functionalities. Tasks may run once or repeatedly. If they run repeatedly the period is based on their execution frequency.

TaskManager

The TaskManager controls the scheduling and execution of all tasks in the framework. It maintains the Framework's internal clock measured in *ticks*. Tasks are scheduled for execution in accordance with this clock.

Task Queue

A queue maintained by the TaskManager, as a schedule of the execution of tasks. Each position in the queue represents a tick and can contain 0 or more tasks. The TaskManager advances one position at a time, executing all tasks scheduled at that position (for that tick). It does not proceed to the next position until all tasks scheduled for the current position have finished running.

TaskSpawner

A TaskSpawner is an assistant to the TaskManager, which manages tasks' status and determines what to do with a task after each time it runs.

Tick

Each position in the task queue represents a discrete instant in the application's time. These positions are termed 'ticks'. All tasks scheduled for a particular tick are executed concurrently and the TaskManager does not advance to the next tick until all tasks scheduled for the current tick finish.

Total Activation

An activation value that represents the total salience of an element. It is a function of both current activation and base-level activation.

Transient Episodic Memory

Transient Episodic Memory is a module in the LIDA Model. It is the episodic memory that decays after a few hours or up to a day. Declarative memories are formed offline from transient episodic memories by a process called consolidation.

Workspace

A module that receives and stores content from several modules including current percepts from PAM, recent local associations from Episodic Memory, and the recent contents of consciousness from the Global Workspace. Structure-building codelets operate on Workspace content, integrating structures, creating new Nodes, and building structures representing higher-level phenomena (events, plans, imaginations, etc.).
